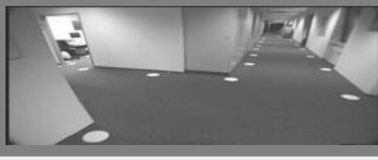

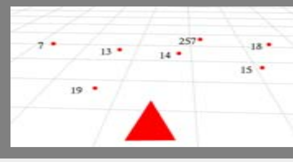

Implicit Data Association from Spectrally-Clustered Local Matches

Edwin Olson
eolson@mit.edu

~~Computer Science and Artificial Intelligence Laboratory~~
~~Massachusetts Institute of Technology~~

Computer Science and Engineering
University of Michigan, Ann Arbor

Motivation

- Data Association is tricky
  - Perceptual Ambiguity
    - *Different* environments look the *same*
    - The *same* environment can look *different*
  - Positional uncertainty exacerbates the problem
    - More potentially similar looking environments
- "Classic" approach
  - Explicit global data association
  - Constellations of features to increase robustness
    - Exponentially-increasing search space
    - Only practical due to efficient search heuristics

Neira & Tardos, 2001
Bailey, 2002

Finding another way…

- Global Place recognition (Haeberlen et al., 2004; Cummins & Newman, 2007; Ho & Newman, 2007)
  - "Never mind the prior, I see something globally unambiguous. I know where I am."
- Local matches (Gutmann & Konolige 2000)
  - "These two local maps look similar. Maybe they're the same place?"
  - Ignore the global context
    - Makes problem smaller → faster to solve
    - False positives!
      - *How large do the local matches need to be?*
      - *How can we detect the "picket fence" problem?*
      - *Larger local matches increase computational costs?*

Our Central Message

- Local matches *are* a reliable and efficient way of performing global loop closing!
- We describe sufficient conditions when a local match is a global match
- Efficiently compute local matches from *very fast* pose-to-pose matches
  - *Exploit* false positives of pose-to-pose errors in order to prevent false positives due to environmental ambiguities
  - Exceptionally few (if any) false-positives end-to-end
  - Never explicitly construct large local maps for matching
    - Avoid computational costs

Agenda

- When are local matches also global matches?
  - Two sufficient criteria
- Building local matches
  - Pose-to-Pose matches
  - Combining pose-to-pose matches into *local* matches
  - Testing the sufficient criteria
- Results
  - DLR Circles Dataset

Agenda

- When are local matches also global matches?
  - Two sufficient criteria
- Building local matches
  - Pose-to-Pose matches
  - Combining pose-to-pose matches into *local* matches
  - Testing the sufficient criteria
- Results
  - DLR Circles Dataset

Global Match Criteria – Global Sufficiency

- A *must* be somewhere inside the uncertainty ellipse
- If uncertainty is large, then area A and B might not be the same place
  - Area A might be *somewhere else* inside uncertainty ellipse

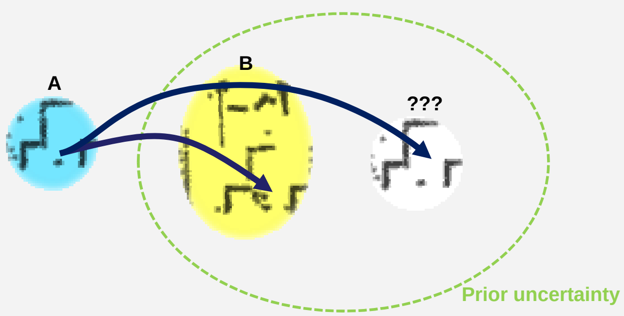

Global Match Criteria – Global Sufficiency

- If match is large, then A must be B
  - No other place to "put" A

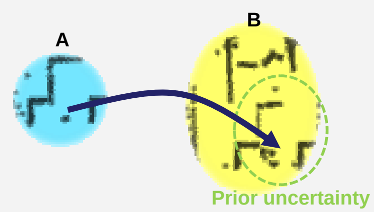

Global Match Criteria – Global Sufficiency

- See??? I told you so!

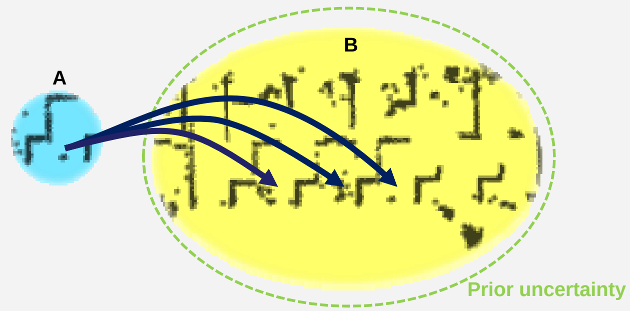

Global Match Criteria – Local Unambiguity

- Second criterion: local unambiguity
  - “Picket Fence Problem”: large overlapping matches
  - Ambiguous!

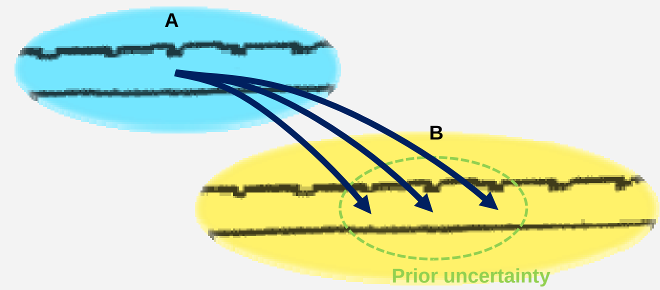

Global Match Criteria: Summary

1. Global Sufficiency: There is no disjoint match (“It’s not somewhere else entirely”)
2. Local unambiguity: There are no overlapping matches (“It’s either here or somewhere else entirely”)

If both are satisfied, A must match B.

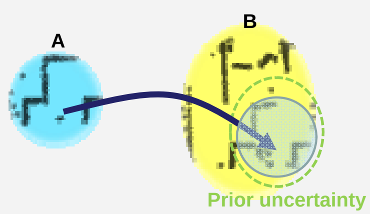

Agenda

- When are local matches also global matches?
  - Two sufficient criteria
- Building local matches
  - Pose-to-Pose matches
  - Combining pose-to-pose matches into *local* matches
  - Testing the sufficient criteria
- Results
  - DLR Circles Dataset

Our Approach: Big Picture

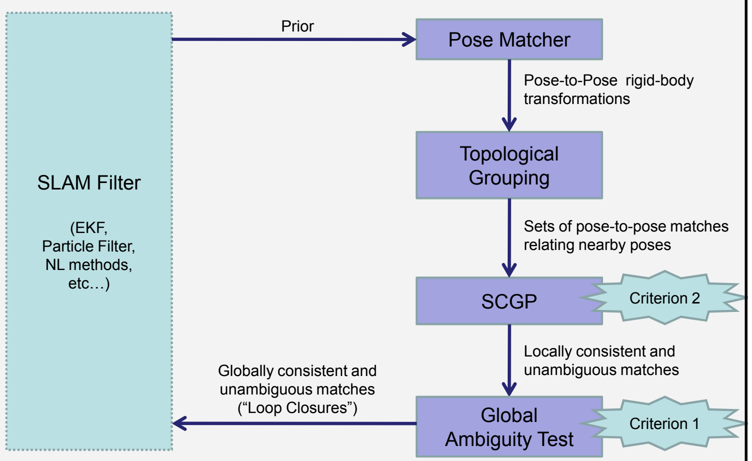

Our Approach: Big Picture

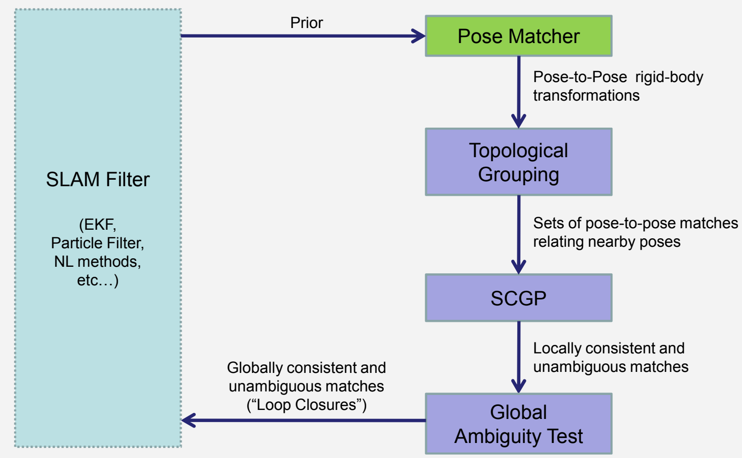

DLR Circles Dataset Preview

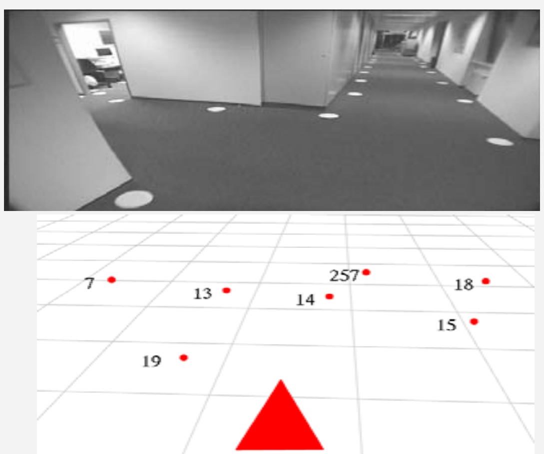

Pose-to-Pose Matching: Plausibility

- Do two poses possibly have overlapping sensor observations?

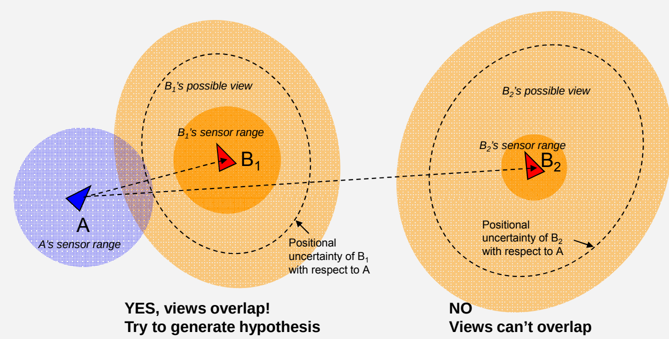

Pose-to-Pose Matching: Plausibility

- Find all poses that might overlap our current pose **a**
  - i.e., compute $P(x_b \mid x_a)$ for all b
- Straightforward in EKF/EIF implementations
  - Manipulate the covariance matrix
- But exactness isn't critical:
  - Can use Dijkstra projection — Bosse , 2004
    - Fast, conservative estimates
  - Gives us flexibility in choosing SLAM implementation
    - e.g., Nonlinear optimization methods

Pose-to-Pose Matching

- Find pose-to-pose correspondences
  - Estimate sensor overlap based on assumed correspondences
  - RANSAC
    - Assume two correspondences, compute optimal rigid-body transformation — Horn, 1987
    - Consensus = # of matched points - # of missing points inside sensor overlap region
    - Best rigid-body transformation is returned
- Critically,
  - this process will produce incorrect rigid-body transformations when the environment is ambiguous

Pose-to-Pose matching

- Example (real data from DLR circles)

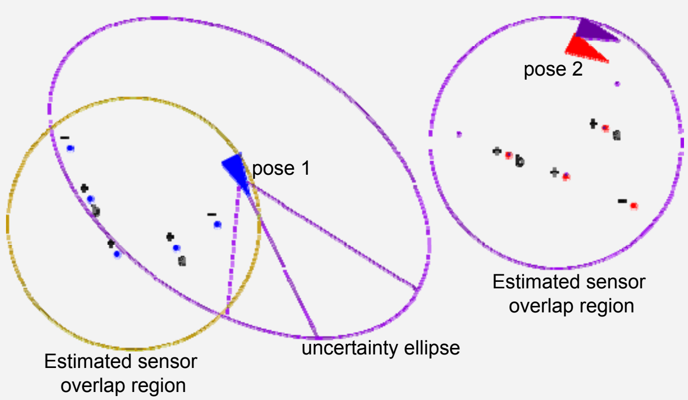

Big Picture

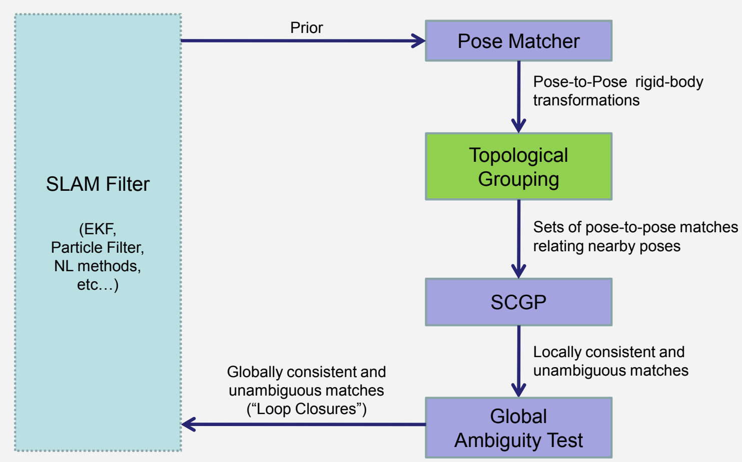

Topological Grouping

- Group together topologically-related pose-to-pose matches to form local matches
  Olson 2008
- Each group asks a "topological question"
  - do two local maps match?

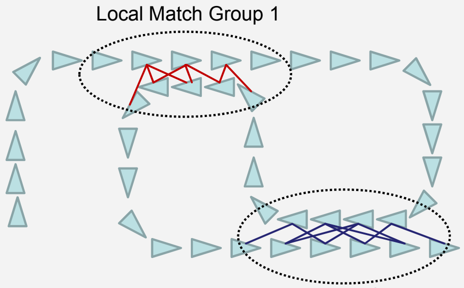

Our Approach: Big Picture

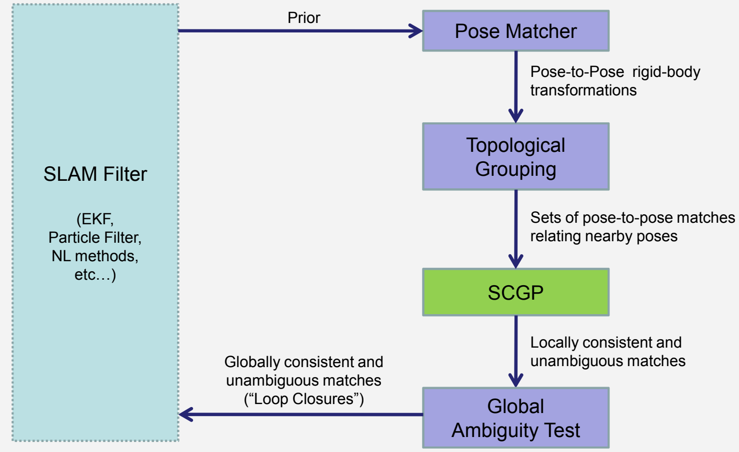

Local Unambiguous Matches

- Our goal:

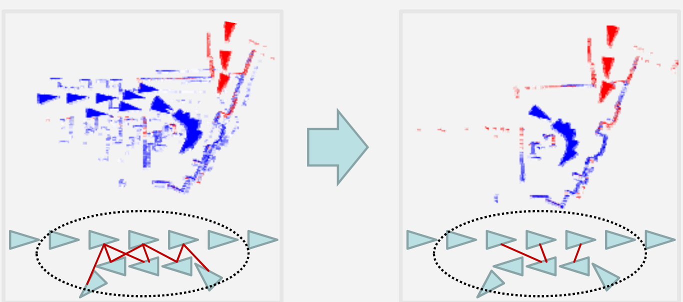

Unfiltered Local Match (set of pose-to-pose matches)

Locally consistent and unambiguous local match (set of pose-to-pose matches)

Locally-Consistent Matches

- Idea:
  - Correct pose-to-pose hypotheses *must* agree with each other
    - (There is only one "truth")
  - Incorrect pose-to-pose hypotheses tend to disagree with each other
    - (There are many ways to be wrong)
  - Find self-consistent subset of hypotheses
    Olson 2005
    - If there are multiple subsets that are self-consistent, then we might have a "picket fence"!

How can we tell if two hypotheses agree?

- Consider two hypotheses **i** and **j** in the set:

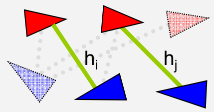

- Form a loop
  - Add two additional edges from our prior

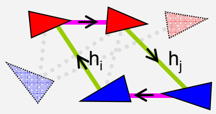

Rigid-body transformation around loop should be the identity matrix

Our Method

- Form pair-wise consistency matrix **A**

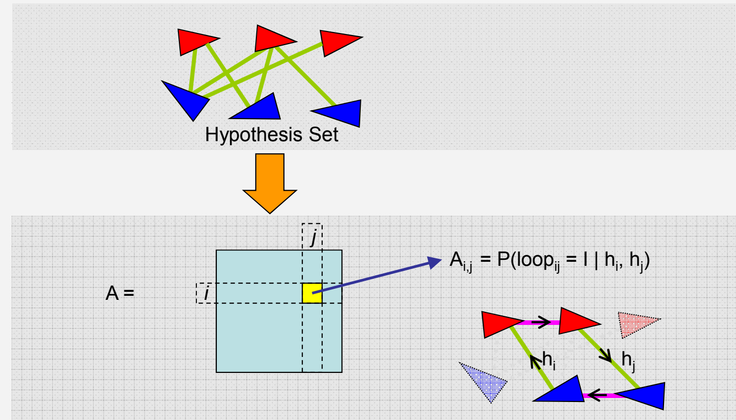

Single Cluster Graph Partitioning [Olson2005]

- Our goal: find the best *indicator vector*
  - Indicator vector represents a subset of the hypotheses

Indicator vector v

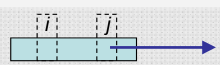

$v_i = 1$ if h_i is correct, 0 if h_i is incorrect

- ***Idea:*** Identify the subset of hypotheses that is maximally self-consistent
  - What subset **v** has the greatest average pair-wise consistency, λ?

$$\lambda = \frac{v^T A v}{v^T v}$$

v^TAv: Sum of all pair-wise consistencies between hypotheses in v

v^Tv: Number of hypotheses in v

  - aka Densest Subgraph Problem

Gallo et al 1989

Consistent Local Matches

- We want to maximize λ by finding a good v:

$$\lambda = \frac{v^T A v}{v^T v}$$

- The best solutions maximize the equation:

$$Av = \lambda v$$

- Critical results:
  - Subset v_1 is maximally self-consistent subset
  - If λ_1 / λ_2 is large (>2) then v_1 is locally unambiguous

Consistent Local Matches

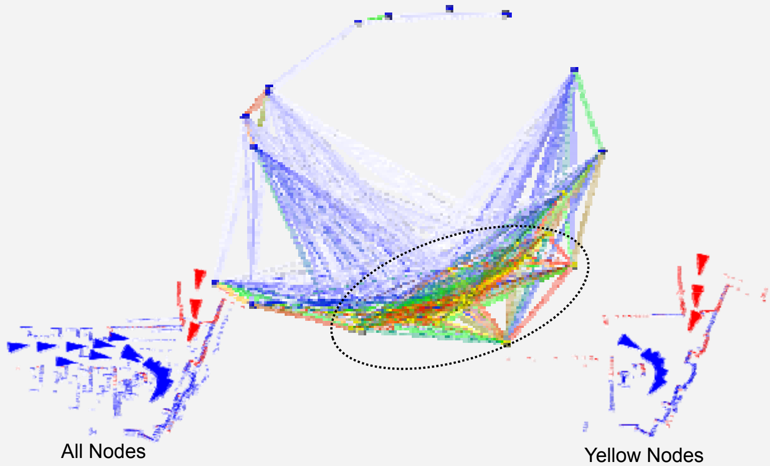

Our Approach: Big Picture

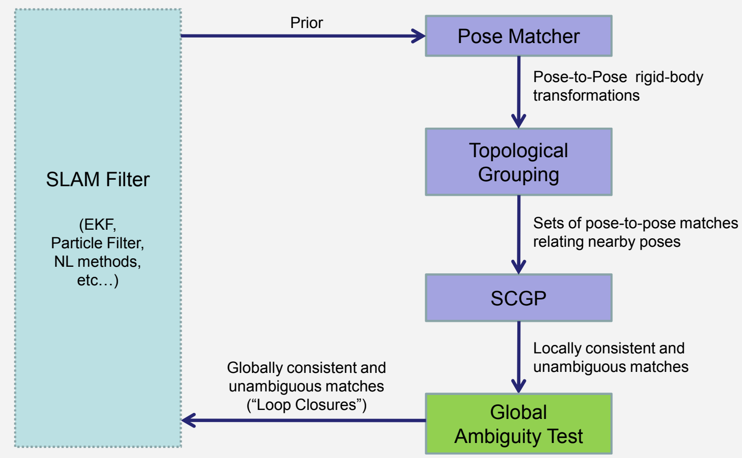

Global Consistency

- Correct method:
  - Can two copies of A be arranged so that they both fit inside the covariance ellipse?
- Our Practical approximation:
  - Is the bounding dimension of A at least half the length of the dominant axis of the covariance ellipse?
  - Potential failure mode: very narrow local matches

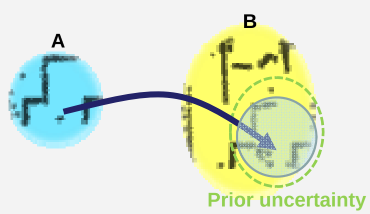

Agenda

- When are local matches also global matches?
  - Two sufficient criteria
- Building local matches
  - Pose-to-Pose matches
  - Combining pose-to-pose matches into *local* matches
  - Testing the sufficient criteria
- Results
  - DLR Circles Dataset

DLR Results

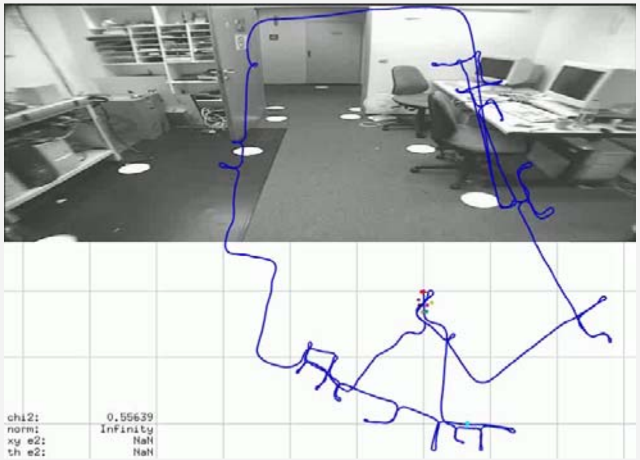

DLR Results

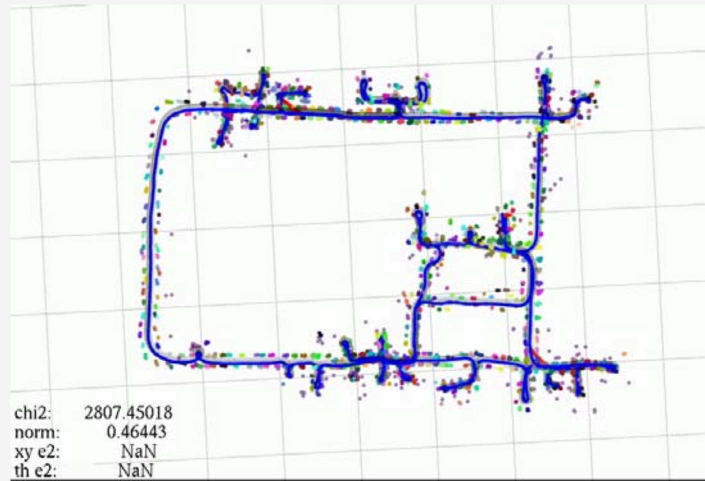

DLR Results

- Annotate each pose-to-pose hypothesis with whether the data-associations match ground-truth
  - Incorrect matches rejected (zero false positive rate)
  - Fairly high false-negative rate
    - But observation noise can cause a hypothesis to be bad even if it was derived from correct associations

Hypothesis Outcome	Good Assoc.	Bad Assoc.
Accepted Hypotheses	2043	0
Rejected (small set)	714	32
Rejected (ambiguous)	1168	105
Rejected (inconsistent)	544	38

- Total CPU time (pose-to-pose matching, grouping, SCGP, graph optimization): 199s

Additional Results

- Also successfully applied to:
  - Standard LIDAR benchmark datasets
  - Bearing-only vision features with highly ambiguous descriptors

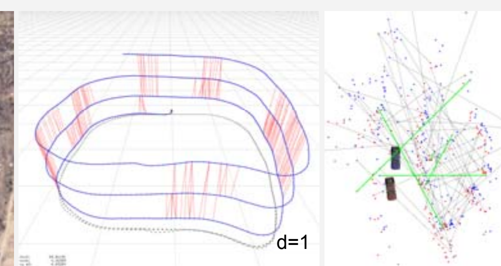

Conclusions

- Local Matching can be used to establish global matches in a principled fashion
- Local Matches cheaply assembled from trivially-computed pose-to-pose matches
- Account for local ambiguity: “picket fence” by using SCGP’s confidence metric
- Account for positional uncertainty: more uncertainty requires more evidence (and the converse too!)
- Friendly to fast SLAM methods
  - Don’t need full covariance